

Thought leadership

GPU infrastructure – financing and contracting for AI compute capacity

July 2026



GPU infrastructure – financing and contracting for AI compute capacity

Key takeaways

- 1 AI compute capacity is driving an infrastructure supercycle in which value creation depends on coordinated investment and alignment across the entire AI compute stack, from chips and AI-ready data centres to power, cooling and connectivity and on how those components operate together within the wider AI compute ecosystem.
- 2 As AI infrastructure evolves as an investable asset class, bankability is increasingly driven by contractual frameworks, revenue certainty and ecosystem alignment and risk allocation, rather than the underlying asset alone.
- 3 Sovereignty, geopolitics and evolving regulatory requirements are becoming core design and financing considerations that must be addressed from the outset to secure investment and long-term viability.

The rapid build-out of AI infrastructure and growing demand for AI compute are reshaping how assets in this space are financed, and how deals are structured and contracted for. As sectors increasingly depend on AI and other data-intensive processes, a robust AI compute stack (spanning data centres, high-performance chips such as GPUs and other AI accelerators, software orchestration, connectivity), and access to the power and cooling required to operate these systems at scale is fundamental.

A key question is how value is distributed across the assets being financed, and how the ecosystem must operate in a co-ordinated way to generate strong returns on investment. The key challenge is no longer simply access to chips; power availability, grid access and water for cooling capacity are also critical factors, increasingly shaping where assets are located, where investment flows and how projects are structured. At the same time, shorter hardware lifecycles and emerging AI-as-a-Service and GPU-as-a-Service models are driving financing structures that look increasingly to contracted revenues, utilisation and risk allocation rather than the underlying asset alone.

In this briefing we explore where capital is being deployed, what makes AI compute assets investable and how contractual frameworks are evolving to support this rapidly developing market.



What is being financed and where is capital going?

We are now entering what has been described as an “AI infrastructure supercycle.” Hyperscaler capital expenditure alone is projected to reach over US\$1 trillion annually by the end of 2027. There is also significant investment across the wider ecosystem, with global data centre investment expected to reach around US\$6–7 trillion by the end of 2030 to support demand for AI and digital infrastructure.

At the hardware layer, semiconductor and AI computing infrastructure account for more than US\$3 trillion of that total, reflecting the cost and importance of GPUs, custom chips and other high-performance systems. Investment is also flowing into energy and supporting infrastructure. With global data centre electricity demand already at around 415 TWh and expected to more than double by 2030, substantial investment is being directed towards power generation, grid capacity and cooling.

“Against that backdrop, a key question is where value sits within the AI compute ecosystem and how value is realised and sustained over the long-term,” says Charlotte Walker-Osborn, a Knowledge Director in Clifford Chance’s global Technology Group. She continues: “Value creation depends on access to and availability of assets across the entire AI compute stack (including real estate, critical build materials, advanced semiconductor chips, networking, and power and cooling infrastructure), and on how those components operate together within the wider AI compute ecosystem. Long-term value is increasingly shaped by the alignment of infrastructure providers, technology suppliers, capital providers, operators, customers and wider supply chain participants, together with financing structures, contractual frameworks and revenue models that are designed to support bankability and sustainable returns from the outset. Overlaying all of this is a rapidly evolving regulatory environment that can directly influence how these assets are structured, financed and monetised.”

Matt Dunn, a London-based Partner, has seen a heavy focus on financing of physical infrastructure over the past two years. “More recently, however, financing has moved beyond single asset real estate deals, particularly in Western Europe, to include sites where capacity is not yet contracted for. Given permitting constraints and long lead times, lenders are increasingly financing powered land without a letter of intent or contract in place. The next stage is portfolio-level financing. Here, the focus is on more flexible structures. US developers are looking to retain flexibility, recognising that moving data centres within corporate structures can trigger tax and regulatory issues. The aim is to keep assets static while allowing ownership and financing arrangements to move around them,” he says.

“Financing has moved beyond single asset real estate deals, particularly in Western Europe, to include sites where capacity is not yet contracted for. Given permitting constraints and long lead times, lenders are increasingly financing powered land without a letter of intent or contract in place.”



Matt Dunn
Partner, London

There is a growing focus on AI infrastructure. This goes beyond campus financing to models where campus-level funding sits alongside GPU-as-a-Service and other AI compute structures (such as AI cloud), sometimes with different borrowers on the same site. Financing individual components of the stack, particularly GPUs and other AI-ready TPUs and AI accelerators, is also increasingly important. This can span vendor financing, equipment leasing, or full GPU-as-a-Service models or other wider AI compute models. GPUs and other high-performance AI chips, are expensive to acquire and operate. As a result, operators are increasingly leasing access to GPUs and other AI accelerators housed in data centres. “What is emerging is a financing model built around those GPU-as-a-Service lease cash flows. Data centres can secure capital upfront by monetising future income streams, enabling them to acquire GPUs in the first place,” says Simon Connor, a Clifford Chance Partner based in London.

In effect, this functions as a structured working capital solution. Data centres acquire GPUs and other AI compute hardware and lease that capacity simultaneously whilst raising financing against contracted revenues. For this to work, every part of the structure must be financeable from the outset. That includes manufacturer contracts, data centre acquisition agreements and lease arrangements with offtakers. “Financing is no longer a secondary step, it is fundamental to how these platforms are designed,” says Matt Dunn.

Across the market, investor approaches vary. Some are vertically integrated, investing across talent, chips and data centres to retain control. Others remain focused on specific parts of the value chain. In both cases, co-ordination across the ecosystem is becoming increasingly important for full value to be achieved.

The rise of circular capital

A form of circular economy is clearly emerging around GPUs and AI compute generally. A prominent example is Nvidia's investment in OpenAI, with OpenAI in turn leasing Nvidia chips for compute capacity. Similarly, Nvidia's investment in AI cloud provider, Lambda, sees Lambda purchasing chips, with Nvidia effectively leasing back usage for its own AI workloads. Beyond Nvidia, similar dynamics are emerging elsewhere. For example, SoftBank is developing its own chips while also investing in OpenAI. The result is a highly interconnected, strategically layered market.

"Less obvious are the dynamics involving neocloud providers," says Gianluca Bacchiocchi, a Clifford Chance Partner based in New York. "For example, companies such as Fluidstack have entered into arrangements with Google backstopping their lease payment obligations. These structures may be designed not only to support financing but also to incentivise neoclouds to use Google's chips or other services," he says. "There is a clear trend of hyperscalers and chip providers becoming more deeply embedded across the entire value chain, from development through to offtake and ongoing operations."

"There is a clear trend of hyperscalers and chip providers becoming more deeply embedded across the entire value chain, from development through to offtake and ongoing operations."



Gianluca Bacchiocchi
Partner, New York

Regional divergence

There are clear regional differences around the world but also signs of convergence. In structured finance, particularly ABS, US clients often push for a single master trust that allows assets to rotate across jurisdictions, but in Europe differences in capital markets frameworks, security structures and rating agency expectations mean a fully unified approach is not always feasible. Even so, the direction of travel is towards greater alignment, with the potential for multi-jurisdictional assets to sit within a single securitisation pool.

"Regionally, the Nordics continue to stand out," says Matt Dunn. "Political stability, abundant low-cost renewable power and aligned permitting regimes have made the region a focal point for both development and financing innovation." This is reflected in sustained M&A activity, including large portfolio transactions, as well as early adoption of new financing structures such as private credit.

There are also growing cross-border capital flows into these markets. Investors from outside Europe are increasingly active, in some cases driven by familiarity with specific offtakers. Transactions such as recent financings in Finland illustrate how global capital is following demand where the fundamentals are strongest.

At the same time, investor attitudes within Europe have shifted. A few years ago, portfolio financings often included limits on exposure to southern European jurisdictions. That has largely fallen away. Hyperscaler demand is driving expansion into markets such as Athens, Madrid and Barcelona, in addition to Warsaw, regardless of historical preferences. In practice, customer requirements are reshaping geographic strategy. If a major offtaker such as Microsoft needs capacity in a given location, operators and investors are following.

Power, grid access and cooling are also becoming decisive location factors. In some regions, access to renewable power, water for cooling and grid connectivity are as important as real estate or latency considerations. This is changing where assets are located and how they are financed.

The impact of digital sovereignty on deal flow and investment

“Sovereignty is not new, but the rise of sovereign AI has materially shifted the market. It is no longer a policy overlay applied to existing infrastructure. It has become a core requirement that shapes how AI infrastructure is designed from the outset,” says Patrice Navarro, a Clifford Chance Partner based in Paris.

Historically, sovereignty focused on data location, data storage and data flows. That approach is no longer sufficient. Today, requirements extend to control of infrastructure, access rights, equipment, remote operations and who ultimately exercises effective control. The question is no longer just where data sits, but who controls the infrastructure and systems and who can access them.

This is not limited to Europe. Similar dynamics are emerging globally, including in India and the US, although approaches differ. Across markets, sovereignty is moving from a “nice to have” to a prerequisite in many countries, particularly where public funding or strategic offtake is involved. This is a trend which is likely to accelerate, although some countries are pushing against it.

Sovereignty must now be considered from day one and designed in where applicable. Rather than retrofitting existing assets, developers are structuring ownership, governance and operational models upfront to meet evolving requirements. This is proving more effective than earlier efforts to impose sovereignty on legacy infrastructure.

For investors and lenders, this adds another layer of diligence. If sovereignty requirements are not met, access to customers, revenues or public funding may be constrained. As a result, understanding how assets are structured and controlled is becoming as important as the underlying commercial model.

Walker-Osborn says: “This diligence increasingly needs to sit alongside foreign direct investment, export control, trade and technology transfer analysis. The regulatory landscape for advanced chips, data centre assets, AI systems and related technology is evolving quickly, and cross-border structures need to be capable of adapting to those changes.”

No single model for deal structures and financing

We are already seeing multiple models operating on the same site; for example, campus-level financing alongside GPU-as-a-Service and financing for AI compute. That creates complexity around contracts, customer agreements and how different financing layers interact.


“One trend to watch in Europe is the co-location of power generation and data centres and the move to finance both together. This is already established in the US and is starting to take hold in markets facing power constraints,” says Matt Dunn.

Ireland is a clear example. Restrictions now require larger developments to generate their own renewable power and to provide dispatchable generation and/or storage capacity that matches their maximum import capacity. We have also seen examples of off-grid, on-site power generation and battery energy storage systems being installed in Irish data centre developments which, in some cases, avoid relying on the grid entirely. That makes standalone financing more difficult and is pushing the local market towards integrated structures that combine generation and data centre assets. Similar pressures are emerging in other European hubs, including Amsterdam.

Public funding is beginning to shape structure, at least in some geographies. EU-backed AI infrastructure initiatives, including proposed gigafactories, are expected to blend public and private capital, with around a third of funding coming from EU agencies and the remainder from private sources. Delivering this will require new models, potentially involving development finance institutions or public-private partnerships.

These structures may also raise additional legal and structuring issues, including state aid/subsidy control, procurement, governance, reporting and eligibility requirements attached to public or multilateral funding.

Overall, financing is becoming more complex and more bespoke. Structures are being driven not just by asset type, but by power constraints, policy considerations and the growing role of public capital.



What makes AI compute assets investable?

“From a legal standpoint, with hyperscalers, you are typically dealing with investment grade counterparties and long-term contracts well in excess of the financing, limited termination rights, as well as the wider non-contractual environment in terms of stickiness and availability zones. You are not taking construction risk, and performance risk is relatively contained,” says Matt Dunn.

However, the picture is slightly different for GPU-as-a-Service and AI-as-a-Service models – there is less track record with counterparties and a different risk profile, including installation and technology risk.


That distinction matters because the financing analysis may shift from underwriting a relatively stable real estate or data centre asset to underwriting a contract, cash flow and service performance model. In many cases, lenders are asking whether they are lending against the hardware or other asset, the customer contract or the wider ecosystem of supply, support and offtake arrangements.

“What is interesting from a bankability perspective is how closely we are seeing partnerships between corporates and banks in setting up these structures. Contractual terms are effectively being diligenced and amended in real time, with banks supporting and shaping arrangements in parallel so they can ultimately lend against them,” says Simon Connor. “That dynamic makes these structures more investable, because we are not diligencing something inflexible. Instead, all parties recognise that the manufacturer cannot sell, the data centre cannot buy and the offtaker cannot lease unless the contracts are drafted in a way that works for everyone. It is a live process, but a constructive one that requires engagement and some original thinking,” he adds.

In the US, there is a significant amount of activity in the capital markets, with a number of interesting deals taking place. In 2025, Meta and Blue Owl formed a joint venture around a 2GW project in Louisiana, raising over US\$27 billion of debt. The SPV created by Blue Owl issued the 144A/Reg S debt into the capital markets. “From Meta’s perspective, they were not issuing the debt, and neither was the joint venture itself. And this was not a traditional financing by any stretch, because it did not have the typical collateral package you would expect in a project finance transaction for a data centre. The structure allowed Meta to deconsolidate the asset, achieve equity treatment from both an accounting and ratings perspective, and – importantly – through the contractual framework have the lease treated as an operating lease rather than a finance lease. That was based on a four year lease with four year renewal terms and a guaranteed residual value,” says Gianluca Bacchiocchi. “This off balance sheet structure took the market by surprise, and others are now looking at it. We have not yet seen a notable direct repeat in the data centre space, but we know more are coming,” he adds.

Separately, capital markets deals are moving in three broad directions. First, high-yield transactions. Secondly, investment grade deals where the bond rating is still materially below the credit quality of the hyperscaler offtaker; for example, a double A tenant supporting a triple B bond. And thirdly, construction risk deals where there is a much closer alignment between the bond rating and the underlying credit quality of the offtaker.

There is no single reason why one structure is chosen over another, but it is clear that sponsors are exploring all available options. “We are also seeing developers going to the capital markets earlier in the life cycle during construction. This has not yet fully played out in Europe or Asia, but we do expect this model to arrive in other markets over time,” says Bacchiocchi.



Structuring for performance risk

“In a typical receivables deal, companies want little or no performance risk – they want goods delivered and services performed. And this is where the market historically has sat. Here, however, you’re dealing with a live service arrangement with performance being a key requirement and therefore a deal where, if service levels drop, the offtaker may have the right to not pay. The core question is how lenders get comfortable with this dynamic,” says Simon Connor.

There are three core elements. First, structuring, diligencing and auditing performance risk. Secondly, the multiple recourse routes. And finally, the business generative nature of the model.

On the first, there is a much more active focus on historic performance for the contract(s) – service levels, delivery metrics – but also on building in real time monitoring. That includes trigger events, termination events, stop purchase rights and acceleration mechanisms, all linked to how the contract is performing on a live basis. The idea is to adjust the lender’s risk exposure dynamically as performance evolves.

“This is a shift from traditional receivables deals, where you rely on historic data – a known dilution rate, for example – and bake that into reserves or advance rates. Here, you may not have that history. But with strong live reporting, you can at least track performance as it develops. Over time, you could see more dynamic approaches – such as dilution reserves that move up or down depending on actual performance, penalising underperformance and rewarding consistency,” says Connor.

On recourse, while performance risk is present, it is not new. We see it in data centre securitisations today, where termination rights exist but are managed through structure, asset quality and alignment of interests. The same applies here.

“What is more interesting is the additional recourse being layered in. Manufacturers are, in some cases, providing quasi vendor financing through backstops on GPU-as-a-Service lease payments – effectively guaranteeing a portion of the cash flow that should be expected to be received from offtakers. That is a meaningful recourse route for financiers and is already being relied upon in some structures,” Connor says.

Over time, you could see more dynamic approaches – such as dilution reserves that move up or down depending on actual performance, penalising underperformance and rewarding consistency.”



Simon Connor
Partner, London

On top of that, you may see security over hard assets – perhaps more of a ‘belt and braces’ approach in the early stages of the market – as well as corporate guarantees. Corporate entities are effectively standing behind performance obligations, providing additional comfort that lease contracts will be serviced as expected.

Taken together, these features go a long way towards explaining how lenders get comfortable with performance risk and are willing to finance against these assets.

Finally, and importantly, this is a business generative model. The financing is not incidental – it is fundamental. It supports the acquisition of the AI compute assets and underpins the leasing model itself. That interplay between financing and the underlying commercial arrangement is what makes this such an active and evolving space.

Managing risk through contract design

“A contract can look strong on paper but deliver materially less in practice,” says Patrice Navarro. “Service credits and penalty mechanics can quietly erode what looks on the face of it like a robust take-or-pay or revenue commitment, so contracts need to be reviewed clause by clause. Most contracts in this space are novel to some degree. These types of deals did not exist 18 to 24 months ago, so there is no real precedent, no established market standard. You cannot rely on history.”

As a result, the pre-signing review needs to be granular. You are testing accountability provision by provision because, once the contract is signed, each step becomes a building block for the next. If you miss something early, it is difficult to fix later. You also need to look closely at the term of the contract and ensure it works for lenders and the financing structure. The level of detail required is significant because diligence after the fact is too late.

“And that is where diligence comes into sharp focus. Some of the more traditional points still really matter – assignment rights, step in rights, the ability to replace a party if something in the ecosystem is not working,” he says.

“A contract can look strong on paper but deliver materially less in practice. Service credits and penalty mechanics can quietly erode what looks on the face of it like a robust take-or-pay or revenue commitment, so contracts need to be reviewed clause by clause.”



Patrice Navarro
Partner, Paris

Other points are increasingly central to bankability. These include whether there is any termination for convenience, how performance and availability regimes operate, when service credits could tip into termination risk, whether minimum commitments or take-or-pay obligations are robust enough to support financing, and how far key supply chain dependencies are flowed down into customer contracts.

This is particularly important where the relevant asset may depreciate quickly, utilisation may fluctuate and value is driven more by contracted revenues than by residual hardware value. In those circumstances, the customer contract cannot be diligenced in isolation. It needs to be reviewed alongside upstream supply contracts with chip suppliers, cloud providers, software vendors, data centre operators, power providers and other infrastructure partners and suppliers.

As a result, the emerging best practice is to align the commercial model, customer contract, supply chain arrangements and financing structure from day one. Retrofitting bankability after contracts have been signed is materially harder.


What makes AI infrastructure bankable?

The rise of AI brings significant opportunities for financing. What is clear from our work with clients is that bankability can be driven just as much by contract structure as by the underlying asset, and that success depends on aligning customer contracts, supply chain arrangements, and financing structures early, rather than trying to retrofit them later.

Additionally, these areas are becoming increasingly important:

- First, diligence is expanding beyond traditional technical and commercial analysis to include foreign direct investment, export controls, trade restrictions, technology transfer regimes and policy direction. This reflects the fact that AI compute assets sit at the intersection of critical infrastructure and geopolitics, and that these areas can materially affect the economics of AI compute projects.
- Secondly, financing analysis is evolving. Rather than underwriting stable real estate or data centre assets, lenders are increasingly underwriting contract performance, cashflow resilience and ecosystem risk. This requires a different risk appetite and more dynamic structuring over the life cycle of data centre assets.
- And third, sovereignty is becoming a structural design input rather than a compliance overlay. Like the financing structure, where sovereignty is applicable, it needs to be built in from day one. Retrofitting it onto existing assets has historically failed, and the market is increasingly treating sovereignty by design as a precondition to investability where public funding or strategic offtake is in scope.

In this environment, an end to end business case – combining legal, technical, commercial and regulatory analysis – is crucial for successful execution.



Contacts



Gianluca Bacchiocchi
Partner, New York

gianluca.bacchiocchi@cliffordchance.com
+1 212 878 3256



Matt Dunn
Partner, London

matt.dunn@cliffordchance.com
+44 207006 4890



Charlotte Walker-Osborn
Knowledge Director – Tech Group
(UK Lead), London

charlotte.walker-osborn@cliffordchance.com
+44 207006 2662



Simon Connor
Partner, London

simon.connor@cliffordchance.com
+44 207006 5304



Patrice Navarro
Partner, Paris

patrice.navarro@cliffordchance.com
+33 1 44 05 53 71

This publication does not necessarily deal with every important topic or cover every aspect of the topics with which it deals. It is not designed to provide legal or other advice.

cliffordchance.com

Clifford Chance, 10 Upper Bank Street,
London, E14 5JJ

© Clifford Chance 2026

Clifford Chance LLP is a limited liability partnership registered in England and Wales under number OC323571

Registered office: 10 Upper Bank Street,
London, E14 5JJ

We use the word 'partner' to refer to a member of Clifford Chance LLP, or an employee or consultant with equivalent standing and qualifications

If you do not wish to receive further information from Clifford Chance about events or legal developments which we believe may be of interest to you, please either send an email to nomorecontact@cliffordchance.com or by post at Clifford Chance LLP, 10 Upper Bank Street, Canary Wharf, London E14 5JJ

Abu Dhabi • Amsterdam • Barcelona • Beijing • Brussels • Bucharest • Casablanca • Delhi • Dubai • Düsseldorf • Frankfurt • Hong Kong • Houston • Istanbul • London • Luxembourg • Madrid • Milan • Munich • Newcastle • New York • Paris • Perth • Prague • Rome • São Paulo • Shanghai • Singapore • Sydney • Tokyo • Warsaw • Washington, D.C.

*AS&H Clifford Chance, a joint venture entered into by Clifford Chance LLP.

**Clifford Chance has entered into association agreements with Clifford Chance Prague Association SRO in Prague and Clifford Chance Badea SPRL in Bucharest.

Clifford Chance has a best friends relationship with Redcliffe Partners in Ukraine.